Appears in the Proceedings of the 17th International System Safety Conference, August 1999, pp. 598-607

From Bridges and Rockets, Lessons for Software Systems

C. Michael Holloway; NASA Langley Research Center; Hampton, Virginia

Keywords: safety, high integrity systems, software engineering, accident analysis, history

Abstract

Although differences exist between building software systems and building physical structures such as bridges and rockets, enough similarities exist that software engineers can learn lessons from failures in traditional engineering disciplines. This paper draws lessons from two well-known failures—the collapse of the Tacoma Narrows Bridge in 1940 and the destruction of the space shuttle Challenger in 1986—and applies these lessons to software system development. The following specific applications are made: (1) the verification and validation of a software system should not be based on a single method, or a single style of methods; (2) the tendency to embrace the latest fad should be overcome; and (3) the introduction of software control into safety-critical systems should be done cautiously.

Introduction

Articles and books abound warning about the inadequacies of software development practices (refs. 1-9). Often, these inadequacies are attributed primarily to differences between software engineering and traditional engineering disciplines. Differences commonly cited include the following: the inherently discontinuous behavior of software as opposed to the inherently continuous behavior of physical systems, the fact that software does not wear out like physical components, and the relative youth of software engineering as compared to traditional disciplines.

Differences such as these exist, but do not justify the attitude that software is *so* different that nothing can be learned from traditional engineering disciplines. There is much that can be learned, as others have recognized. For example, in a 1994 article Nancy Leveson drew parallels between the early development of high-pressure steam engines and current software engineering. She wrote, "Risk induced by technological innovation existed long before computers; this is not the first time that humans have come up with an extremely useful new technology that is potentially dangerous. We can learn from the past before we repeat the same mistakes" (ref. 10).

This paper is based on a similar premise, but uses a different approach. Instead of looking at the development of a particular technology, we look at two specific failures from two very different technologies: bridges and rockets. Studying failures was chosen because, as Henry Petroski has written, the lessons learned from failures "can do more to advance engineering knowledge than all the successful machines and structures in the world" (ref. 11). Bridges and rockets were chosen for two reasons. First, building bridges is one of the oldest engineering activities, and building rockets is one of the youngest. Second, the collapse of the Tacoma Narrows Bridge in 1940 and the destruction of the space shuttle Challenger in 1986 are two of the most widely known engineering failures of this century.

The discussion of both failures will be necessarily brief and incomplete, and will contribute nothing new to the understanding of either. The paper's contribution is in the direct application of lessons from these failures to software engineering.

The structure of the remainder of the paper is simple. First, the Tacoma Narrows Bridge collapse is described, and four lessons from it are explained. Second, the Challenger accident is described; how this accident reinforces lessons from Tacoma Narrows is explained; and one additional lesson is added. Third, applications of the lessons are made to software systems. Finally, brief concluding remarks are made.

Tacoma Narrows Bridge Failure

Background: The first bridge connecting the Olympic Peninsula with the mainland of Washington was completed in 1940. The suspension bridge was built by the Washington Toll Bridge Authority to provide an alternative to taking ferries across Puget Sound to get to and from the Olympic Peninsula. Constructing the

bridge took only nineteen months, at a cost of $6.4 million, which was financed by a grant from the Public Works Administration and a loan from the Reconstruction Finance Corporation. With a main span of 2800 feet, the bridge was the third longest suspension bridge in the world at that time. Only the George Washington Bridge in New York, and the Golden Gate Bridge in San Francisco were longer (ref. 12 is the source for the material in this section, unless otherwise noted).

The bridge was designed by Leon Moisseiff, who was one of the world's top authorities on bridge design. Moisseiff had been called in to design the bridge after the design proposed by the Washington Department of Highways was rejected as being too expensive. The Department's design called for 25-foot deep stiffening trusses on both sides of the roadway to protect the structure from the strong winds that blew in the Narrows. Projected construction costs were $11 million.

Along with his partner Fred Lienhard, Moisseiff had developed a mathematical theory for calculating load and wind forces for suspension bridges. This theory, called deflection theory, was originally devised by the Austrian Josef Melan, but Moisseiff and Lienhard put it into practice. The underlying idea of the theory was that the "dead load of a suspension structure substantially moderates structural distortions under live load." (ref. 13) Using deflection theory, Moisseiff was able to justify stiffening the bridge with only eight-foot deep plate girders, instead of the 25-foot deep trusses proposed by the Department of Highways. This change was a substantial contributor to the difference in the projected costs of the designs.

Because the amount of traffic over the bridge was expected to be fairly light, the bridge had only two lanes. As a result, the bridge was only 39 feet wide. This was quite narrow, especially in relation to its length. With only the eight-foot deep plate girders providing additional depth, the bridge was also shallow. The resulting silhouette was thought to be both dramatic and graceful.

The narrow, shallow bridge was flexible, more flexible than any other existing suspension bridge. This flexibility was noticed by the builders during construction, and it was also noticed by drivers as soon as the bridge opened to toll-paying traffic on 1 July 1940. At times the bridge undulated so much that drivers would be unable to see cars in front of them as the pavement rose and fell. Some travelers were reported to have even gotten "seasick" when crossing the bridge. The bridge quickly was nicknamed "Galloping Gertie". Traffic on the bridge in its first two weeks was twice what had been expected, perhaps because it attracted not only those who needed to make the crossing, but also the area's roller coaster aficionados.

To reduce the amplitude of the bridge's wave motion, various checking cables and devices were added to it, as they had been to other suspension bridges with greater than expected oscillations. Also, The Washington Toll Bridge Authority contracted with the engineering department at the University of Washington to study how to reduce the bridge's movements. Professor F. B. Farquharson led the investigation, which experimented with a scale model of the bridge in a wind tunnel. Farquharson and his students issued a report suggesting that the bridge could be stabilized by adding additional cables, attaching curved wind deflectors, and drilling holes in the girders to let wind pass through. Disaster struck before the recommendations in the report could be implemented (ref . 14).

The Accident: On 7 November 1940, the clamps holding one of the added checking cables slipped in a wind of about 40 miles per hour. When this happened Galloping Gertie began to move in a new way. Instead of just oscillating up and down as it had before, it started twisting about its centerline. The bridge was closed to traffic, and Professor Farquharson went to observe what was happening.

On the bridge was a logging truck (ref. 14), a car, its owner (a newspaper reporter), and his dog; the driver and passenger of the logging truck had escaped to safety. Farquharson joined the reporter on the heaving deck. Together they tried to get the dog out of the car. As the bridge's motion became increasingly violent, the two men gave up trying to rescue the dog. Instead, they concentrated on rescuing themselves.

The last few minutes of the bridge's demise was captured on film. The resulting footage has probably been seen by just about every engineering student in the last 50 years. On the film, Farquharson and the reporter can be seen trying to make their way to safety. The professor had an easier time of it, because he walked along the centerline of the bridge, which was nearly motionless. The reporter struggled along the edges of the roadway, which was heaving violently. Both made it; the only casualty of the eventual collapse was the dog.

When the amplitude of the undulations in the bridge reached twenty-five feet, the suspender ropes starting tearing, and the deck broke, sending the car and truck into the water. Within 30 minutes, the rest of the deck fell into Puget Sound, leaving only the towers remaining. These towers had been bent out of shape by the violent motion; they were dismantled before a replacement bridge was built.

Investigation: The Federal Works Agency appointed three engineers to investigate the failure: Theodore von Kármán, the director of the Daniel Guggenheim Aeronautical Laboratory at the California Institute of Technology; Glenn B. Woodruff, the engineer of design for the San Francisco-Oakland Bay Bridge; and Othmar Ammann, a world-renown bridge designer. They issued their report less than five months after the collapse occurred.

This report exonerated the bridge designers and engineers saying that "the Tacoma Narrows Bridge was well designed and built to resist safely all static forces, including wind, usually considered in the design of similar structures. … It was not realized that the aerodynamic forces which had proven disastrous in the past to much lighter and shorter flexible suspension bridges would affect a structure of such magnitude as the Tacoma Narrows Bridge" (ref. 15). That is, the engineers had followed the current state of the art. They had used the accepted techniques for taking wind effects into account. As mentioned earlier, these techniques had been developed by Moisseiff himself. It just so happened that these techniques turned out to be flawed.

The report did record, however, that one particular engineer had raised concerns about the design before the bridge was built. Theodore L. Condron was an advisory engineer for the Reconstruction Finance Corporation. His approval of the bridge design was a necessary part of the approval of the loan application to help finance construction. As he studied the plans, Condron became concerned about the narrow width of the bridge relative to the length of its main span. He developed a table (table 1) to compare the ratio of span to width in the proposed design to that of recently completed suspension bridges (ref. 15).

Table 1 - Span to Width Ratios

Bridge	Span(ft)	Width(ft)	Ratio
Delaware River	1,750	89	1:19.7
Ambassador	1,850	59.5	1:31.1
Whitestone	2,300	74	1:31
San Francisco Bay	2,300	66	1:35
Geo. Washington	3,500	106	1:33
Golden Gate	4,200	90	1:46.7
Tacoma Narrows	2,800	39	1:72

This table showed the proposed Tacoma Narrows bridge to be significantly more slender than any other existing suspension bridge. To Condron, this seemed to be going far beyond current experience.

Engineer Condron was sufficiently concerned that he continued to investigate. After hearing that the University of California at Berkeley had conducted some tests on models of suspension bridges, he visited with Professor R. E. Davis in Berkeley. According to Condron, Professor Davis "felt reasonably confident that the lateral deflections of the Tacoma Narrows Bridge as designed and determined by Mr. Moisseiff would be in no way objectionable to users of the bridge" (ref. 15).

Condron found additional support for deflection theory in the written discussion that accompanied the article in which Moisseiff and Lienhard published their theory. The discussion cited the University of California studies as confirming the accuracy of deflection theory. One discussant went so far as to say that Moisseiff and Liendard's analysis was sufficient "to silence all arguments for unnecessary floor widths" (ref. 16). What was lacking in this discussion and in the Berkeley tests, and what would eventually lead to the downfall of deflection theory at Tacoma Narrows, was the

recognition that accounting for lateral deflections alone was not enough: vertical deflections mattered, too.

Because he was the only one who seemed to have doubts about the bridge design, and because the deflection theory of Moisseiff and Lienhard had widespread support among bridge engineers, Condron ultimately acquiesced. He wrote in his final report: "In view of Mr. Moiseiff's recognized ability and reputation, and the many expressions of approval … of his methods of analyses of stresses and deflections in the designs of long span suspension bridges, … I feel we may rely upon his own determination of stresses and deflections" (ref. 15).

His support was not unqualified, however. In his final report he also recommended considering widening the bridge to 52 feet. Had this been done, the width-to-span ratio would have been 1:53.8. The bridge would still have been the narrowest in existence, but less radically so than it turned out to be. According to Petroski (ref. 12), "had Condron's recommendation been followed, it is very possible that the Tacoma Narrows Bridge would have been stiffened enough that, even had it exhibited some degree of flexibility in the wind, that might have been within tolerable limits and thus subsequently correctable, as it was to be in other contemporary bridges."

As we know, Condron's recommendation was not followed, and the bridge collapsed. Some time was to pass before the actual cause of the collapse would be determined (the report from Ammann, Woodruff, and von Kármán left the matter vague). The details of the cause are not important for this paper. What is important is to realize that the theory on which the bridge was designed was flawed because it did not take into account everything that needed to be taken into account. In particular, the dynamic effects of wind load on the bridge were ignored. Reliance on the flawed theory was a significant contributor to the failure.

Relevant Lessons: Some of the lessons of the Tacoma Narrows failure are specific to suspension bridge building. One such lesson, for example, was the need for aerodynamic testing; this testing became a standard procedure in suspension bridge structural analysis in every bridge built afterwards. In addition to such specific lessons, there are at least four lessons with application beyond bridge building. These lessons are explained in the rest of this section.

Lesson 1: Relying heavily on theory, without adequate confirming data, is unwise.

At the time of the design of the Tacoma Narrows Bridge, Moisseiff and most other bridge engineers believed that the accuracy of deflection theory had been adequately confirmed. As mentioned earlier, the results of tests on scale models at the University of California had shown close agreement with the theory's predictions for lateral deflections. Also, several bridges had been designed using the theory, and they were still standing.

As Theodore Condron had suspected, neither the scale model tests, nor the existing bridges truly provided adequate confirming data. The scale model tests were inadequate confirmation because they did not produce any data about vertical deflections. The existing bridges were inadequate confirmation because none of them were nearly as narrow and shallow as the Tacoma Narrows Bridge.

The first real test of the accuracy of deflection theory occurred above the waters of the Puget Sound. When this test failed, the inaccuracy of the theory became apparent. Over time, problems occurred in other bridges that had been designed using deflection theory. Many of them were eventually modified to employ additional means of stiffening.

Lesson 2: Going well beyond existing experience is unwise.

Although many previous suspension bridges had been built, including two with longer spans, the Tacoma Narrows Bridge was unique. As table 1 showed, the span to width ratio of the bridge was 54% greater than that of any contemporary bridge. The Tacoma Narrows Bridge was not a simple extrapolation from existing experience; it was a radical departure from that experience.

Even with deflection theory seeming to justify such a departure, the most prudent action would have been to make small, incremental steps in

narrowing bridge deck widths. This seems especially true when one realizes that the bridge with the next biggest span to width ratio, the Golden Gate Bridge, was at that time showing far greater flexibility than had been calculated. Theodore Condron seems to have been one of the few engineers of the time who had learned this lesson. His advice to widen the bridge to 52 feet was a prudent, incremental step.

Lesson 3: In studying existing experience, more than just the recent past should be included.

The University of Washington's Professor Farquharson continued to study suspension bridges after escaping from Galloping Gertie. In a 1949 report, he gave a historical review of the dynamic behavior of suspension bridges (ref. 17). In this review, he listed ten suspension bridges that were destroyed by wind between 1818 and 1889; nine of these occurred before 1865. He wrote that the failure of the Tacoma Narrows Bridge "came as such a shock to the engineering profession that it is surprising to most to learn that failure under the action of the wind was not without precedent" (ref. 18).

Had Moisseiff and other engineers of his time been aware of this history, and if they had studied the works and writings of such engineers as John Roebling (the designer of the Brooklyn Bridge), they might have been less inclined to dismiss the dynamic effects of wind in the way that they did.

Lesson 4: When safety is concerned, misgivings on the part of competent engineers should be given strong consideration, even if the engineers can not fully substantiate these misgivings.

No one can deny that Theodore Condron's misgivings about the Tacoma Narrows Bridge turned out to be correct, despite his own admission that he could not prove that the design was faulty. As the Challenger accident discussion will show, this was not an isolated case.

Challenger Accident

The discussion of this accident will be briefer than that of the Tacoma Narrows Bridge collapse. In particular, details of the causes of the accident will be discussed only in the context of the relevant lessons. Unless otherwise indicated, the factual information in this section is based on references 19 and 20.

Background: Challenger was one of four vehicles that made up the National Aeronautics and Space Administration's (NASA's) space shuttle fleet; the other three were named Columbia, Discovery, and Atlantis. Before the accident, these four vehicles had flown to space a total of twenty-four times, with Challenger flying the most (nine times) and Atlantis the least (two times).

The basic configuration of all four vehicles was the same. As shown in figure 1, three main components make up the shuttle system: the Orbiter, which houses the crew and payload, and includes the three main engines and the orbital maneuvering system; the External Tank, which holds fuel for the main engines; and two Solid Rocket Boosters (SRBs), which provide about 80% of the thrust for launch. The Solid Rocket Boosters are jettisoned about 2 minutes after liftoff; they are recovered and reused. These Boosters are composed of several sections joined together; one of these joints is labeled in the figure. The External Tank is jettisoned about 8.5 minutes after liftoff; it is not reused.

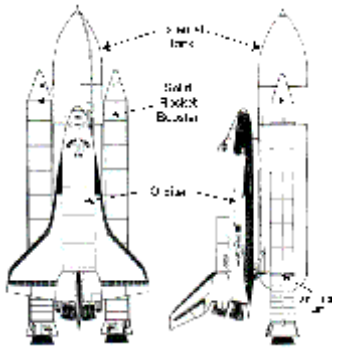

Figure 1 - Space Shuttle Configuration

The Accident: On 28 January 1986, Challenger was scheduled to make its tenth flight into space. The mission had several objectives. These included deploying a Tracking and Data Relay Satellite to support communication with the shuttle and other spacecraft, and deploying the Spartan-Halley satellite, which was designed to study Halley's comet. The part of the mission that made it the subject of more publicity than most previous shuttle missions was that it

carried the first "Teacher-in-Space." New Hampshire schoolteacher Christa McAuliffe was part of the crew. She was scheduled to broadcast a series of lessons to school children across the country during the planned seven day flight.

The launch had originally been scheduled for January 22. Various delays caused successive postponements, until finally Challenger lifted off at 11:38 a.m. on January 28. To spectators watching the launch in person or on TV, everything appeared to be normal. The appearance of a normal flight continued until about 73 seconds after liftoff, when a fireball appeared and the single column of flame and white smoke split into a Y shape, and the orbiter itself seemed to disappear. For nearly an hour afterwards, debris fell into the Atlantic Ocean about 20 miles from the launch site. All seven crew members (commander Francis Scobee; pilot Michael Smith; mission specialists Ellison Onizuka, Ronald McNair, and Judith Resnick; and payload specialists Gregory Jarvis and Christa McAuliffe) died in the accident.

Investigation: A few days after the disaster, President Ronald Reagan established a Presidential Commission to investigate the accident, and charged it with delivering a report to him within 180 days. Former Secretary of State William B. Rogers was appointed as chair of the Commission.

The Commission released their report in June 1986. The Commission "*concluded that the cause of the Challenger accident was the failure of the pressure seal in the aft field joint of the right Solid Rocket Motor*. The failure was due to a faulty design unacceptably sensitive to a number of factors. These factors were the effects of temperature, physical dimensions, the character of materials, the effects of reusability, processing, and the reaction of the joint to dynamic loading." The Commission also concluded, "the decision to launch the Challenger was flawed" (ref. 19, italics in original).

Reinforced Lessons: Three of the four lessons mentioned previously are reinforced by the Challenger accident. One is reinforced by the history of the design of the joint that failed; the other two are reinforced by the events leading up to the decision to launch.

Recall that the second lesson from the Tacoma Narrows Bridge was this: going well beyond existing experience is unwise. At a quick glance, it appears that the designers of the SRB field joints heeded this lesson.

In 1973, NASA Administrator James Fletcher announced that Thiokol Inc. (later to become Morton-Thiokol Inc.) had been selected to design and build the solid fuel rocket motor for the shuttle. In an effort to ensure reliability, while at the same time reducing costs, Thiokol based the design of their segmented booster on that of the Air Force's Titan III rocket. This rocket, which was built by United Technologies, was generally considered as one of the most reliable rockets ever built.

Like the Titan III, Thiokol's design for the field joints had a tang on the rim of one segment slipping into a clevis on the rim of another segment, with the two segments fastened together by pins. While the Titan III had a single O-ring in each joint to seal the joint against the high pressure from the propellant burning inside the booster, Thiokol used two O-rings in the SRB joints (ref. 21). So, Thiokol appeared to be cautiously building on existing experience.

Figure 2 shows an outline of the Titan III and SRB joints next to one another.

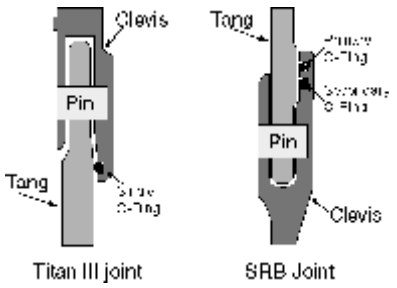

Figure 2 - Joint Comparison

Although this figure does not show all the differences between the joints, it does show an important one: in the Titan III joint, the clevis points downward, but in the SRB joint it points upward. Other differences included the following: to accommodate the second O-ring, the SRB tang was longer than the Titan's, making it more susceptible to bending under combustion pressure; on the Titan the insulation

of the segments fit tightly together, while on the SRB they did not and putty filled the gaps; a single Titan rocket was used only once, but the SRB segments were intended to be reused; and the combustion pressure within the booster was significantly less for the Titan than for the shuttle (ref. 21).

When the details of the joints are compared, it becomes clear that the design was actually just as much a departure from existing experience as the Tacoma Narrows Bridge had been. The resulting failure of the joint reinforces lesson 2.

Lesson 3 (in studying existing experience, more than just the recent past should be included) is also reinforced by the Challenger disaster. In retrospect, it is not difficult to see parallels between some attitudes within the shuttle program before Challenger and some attitudes within the Apollo program before the Apollo 1 fire. The attitude of great confidence in accomplishments and the concern about meeting the planned schedules are especially apparent.
Finally, the Challenger accident also strongly reinforces the fourth lesson: when safety is concerned, misgivings on the part of competent engineers should be given strong consideration, even if the engineers can not fully substantiate these misgivings. Probably everyone who knows anything about the accident knows that on the night before the launch several engineers at Morton-Thiokol argued against launching the next day. In a teleconference with the NASA officials responsible for the SRBs, Thiokol initially recommended against launching until the temperature was above that of the previous coldest launch. After conversations with NASA representatives, and a private caucus among the Thiokol managers and engineers, Thiokol changed their position and recommended launch.

Although other factors may have played a role, one important reason Thiokol managers ended up recommending launch is that their engineers were not able to *prove* by the available data and theories that the launch would be unsafe. The existing data showed that the worst case to date of damage to an O-ring had occurred at the lowest temperature in which a launch had occurred. But the data also showed that the next worst case occurred at a temperature that was one of the highest of all the launches, and that test firings at low temperatures had shown no O-ring damage. The accepted theory at the time also predicted that an O-ring could sustain damage three times worse than any previously experienced and still seal a joint.

Given an equal burden of proof on those who favored launch and those who opposed launch, the decision to launch, although shown by events to have been wrong, was not unreasonable (ref. 22). As lesson 4 implies, the burden of proof ought not to be equal.

A New Lesson: There is at least one more lesson that the Challenger disaster teaches. This lesson is essentially the mirror image of lesson one.

Lesson 5: Relying heavily on data, without an adequate explanatory theory, is unwise.

Many different aspects of the history of the SRB joints could be used to illustrate this lesson, but only one will be discussed here. The booster joints were originally designed with the expectation that the propellant pressure at ignition would cause the inner flanges of the tang and clevis to bend towards each other. This, in turn, would increase the compression on the O-rings and further ensure that they sealed the joint.

In 1977 Thiokol conducted a hydroburst test to assess the strength of the steel case segments. In this test a segment of the booster was filled with oil and put in a chamber. Instruments were attached to a leak check port on a joint to measure the pressure between the two O-rings. The oil was pressurized to about 1.5 times the expected pressure at ignition. The test showed that the steel case was strong enough, but it also showed something completely unexpected. In the first few milliseconds after ignition, the inner flanges of the tang and clevis moved away from each other, thus reducing, not increasing the compression on the O-rings (ref. 21). Figure 3 (ref. 23) illustrates this phenomenon, which was called joint rotation. Notice how the sides of the booster bulge outward, and the joints themselves open up (the effects are exaggerated in the figure so that they can be seen).

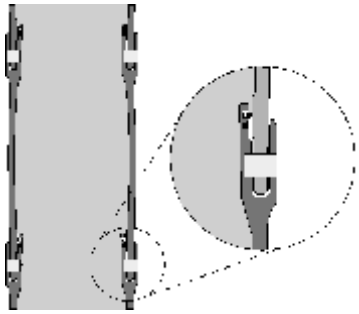

Figure 3 - Joint Rotation (exaggerated)

A 1978 static test firing of a full booster confirmed the existence of joint rotation. Engineers at both Thiokol and NASA were concerned. The two groups disagreed on the actual size of the gap caused by joint rotation. NASA engineers believed the gap was potentially large enough to cause the secondary O-ring to be unable to seal in the event of the primary O-ring failing late in the ignition cycle. The history of the interaction between the two groups is complicated, and not important for this paper. What is important is that eventually both groups were satisfied by the data from various tests and seven static motor firings that the O-rings would seal the joints (ref. 22).

The data convinced them, but no one had a good understanding of exactly why the joints behaved differently than the design predicted they would. The engineers relied on the data without an adequate explanatory theory about why the data was what it was. No such theory ever was developed before the accident (ref 22). Just as relying on theory without sufficient confirming data contributed to the Tacoma Narrows collapse, so too did relying on data with an explanatory theory contribute to the Challenger accident.

Applications to Building Software Systems

Many applications of the five lessons we have just seen can be made to software system development. Only three will be given here.

Application 1: The verification and validation of a software system should not be based on a single method, or a single style of methods. This application is based on a combination of lesson one (relying heavily on theory, without adequate confirming data, is unwise) and its converse, lesson five (relying heavily on data, without an adequate explanatory theory, is unwise).

In the verification and validation of a particular system, this application suggests that neither testing nor analytic techniques should be trusted alone. Testing by itself cannot guarantee the correctness or safety of a system; analytic techniques such as formal modeling are also needed. But, formal modeling should not be used by itself either. No matter how well constructed a formal model may be, rigorous testing of the actual system is still important, especially for validating the accuracy of the assumptions made by the formal model.

Too often, especially at conferences and in the published literature, supporters of testing expend many words showing the limitations of formal methods and supporters of formal methods expend many words showing the limitations of testing. Every testing method has limitations; every formal method has limitations, too. Testers and formalists should be cooperating friends, not competing foes.

Application 2: The tendency to embrace the latest fad should be overcome. Lesson three (in studying existing experience, more than just the recent past should be included) provides the foundation for this application.

Although few software engineers or managers would explicitly claim to be embracing the latest fad, a study of the history of the software discipline shows that it has been characterized by fad-ism.

Famous fads from the past include structured programming, high-level programming languages, artificial intelligence (AI), program verification, and computer-aided software engineering (CASE) tools. Each of these was, at one time, touted by vocal supporters as the solution to the "software crisis." Each of these has contributed in some way to improvements in software. For some, such as structured programming and high-level programming languages, the contributions have been significant, but none of these has come close to delivering the benefits claimed by zealous proponents.

Although Fred Brooks warned over a decade ago against expecting any one particular approach to solve the problems of software development (ref. 24), fad-ism continues unabated. Enthusiasm for object-oriented design and process maturity models remains strong. When this enthusiasm wanes (as it certainly will), architectural design and soft computing seem poised to compete for fad status.

If software practitioners and managers will study history, and learn its lessons, they will stop embracing the latest fads. Instead, they will choose from the wide variety of available techniques those that are most applicable to their particular situation. The quality of software systems will inevitably improve when this happens.

Application 3: The introduction of software control into safety-critical systems should be done cautiously. This application follows straightforwardly from lesson two (going well beyond existing experience is unwise), and is also supported by lesson four (when safety is concerned, misgivings on the part of competent engineers should be given strong consideration, even if the engineers can not fully substantiate these misgivings).

No one intentionally advocates being incautious in using software, but just as the two accidents studied here show, even exceptionally bright people can be self-deceived (ref. 25) about the extent to which their proposals go beyond current experience. Given the complexity of modern software systems, and the tendency of complexity to lead to unexpected accidents (ref. 26), prudence seems to dictate special caution for software systems. Recommendations from software professionals (for one example, see ref. 27), for such caution should be taken seriously, even when these recommendations cannot be fully proven either analytically or empirically.

This does not mean that software should not be used in any safety-critical systems. It already is being used successfully. For example, after studying the design and testing of several Shuttle systems, one of the members of the Rogers Commission expressed greater confidence in the integrity of the software system than in any other system he studied (ref. 23).

Software can be used in safety-critical systems. But its use ought to be guided by successful past experiences, and not by ambitious future dreams. Most children learn to crawl before they walk, and to walk before they run. Software system designers and implementers should do the same.

Concluding Remarks

In 1990 Mary Shaw of the Software Engineering Institute wrote, “Software engineering is not yet a true engineering discipline, but it has the potential to become one” (ref. 28). Her words are no less true today than they were when she wrote them almost a decade ago. Studying established engineering disciplines, and applying the lessons learned in their failures, is one of the ways that the potential of software engineering can be realized. This paper has made a small contribution towards that end.

Although software engineering failures have contributed to loss of life (ref. 29), and to destruction of property (ref. 30), a catastrophe analogous in its public impact to either the Tacoma Narrows Bridge collapse or the Challenge accident has not happened yet. Understanding the fallibility of humans, and knowing a little bit about the history of technology, suggests that such catastrophes are inevitable. Nevertheless, if software engineers and managers are diligent to learn the lessons taught by the past—the past of software engineering, the past of established engineering disciplines, and the past of any other area with relevant lessons—perhaps these catastrophes can be reduced in frequency and in severity. After all, the second bridge over the Tacoma Narrows has been standing for almost 50 years, and the space shuttles have flown nearly 70 safe missions since flights resumed.

References

1. W. Wayt Gibbs. “Software's Chronic Crisis.” *Scientific American* (September 1994): 86-95.

2. Bev Littlewood and Lorenzo Strigini. “The Risks of Software.” *Scientific American* (November 1992): 62-75.

3. Peter G. Neumann. *Computer Related Risks.* New York: ACM Press, 1995.

4. Ivars Peterson. *Fatal Defect: Chasing Killer Computer Bugs.* New York: Times Books, 1995.

5. Evan I. Schwartz. "Trust Me, I'm Your Software." *Discover* (May 1996).

6. Lauren Ruth Wiener. *Digital Woes*. Reading, Massachusetts: Addison-Wesley Publishing Company, 1992.

7. Winston Royce. "Current Problems." In *Aerospace Software Engineering*, edited by Christine Anderson and Merlin Dorman, 5-15. Washington, D.C.: American Institute of Aeronautics and Astronautics, 1991.

8. Christof Ebert. "The Road to Maturity: Navigating Between Craft and Science." *IEEE Software* 14, no. 6 (1997): 77-82.

9. Stuart Shapiro. "Splitting the Difference: The Historical Necessity of Synthesis in Software Engineering." *IEEE Annals of the History of Computing* 19, no. 1 (1997): 20-54.

10. Nancy Leveson. "High Pressure Steam Engines and Computer Software." *IEEE Computer* 27, no. 10 (1994): 65-73.

11. Henry Petroski. *To Engineer is Human: The Role of Failure in Successful Design.* New York: Vintage Books, 1992.

12. Henry Petroski. *Engineers of Dreams: Great Bridge Builders and the Spanning of America*. New York: Alfred A. Knopf, 1995.

13. Robert W. Hadlow. "Historic American Engineering Record: Tacoma Narrows Bridge." NAER No. WA-99. August 1993.

14. Rita Robison. "Tacoma Narrows Bridge Collapse." In *When Technology Fails*, edited by Neil Schlager, 184-190. Detroit: Gale Research, 1994.

15. O. H. Amman, Theodore von Kármán, and Glenn B. Woodruff. *The Failure of the Tacoma Narrows Bridge*. Washington, D.C.: Federal Works Agency, 1941. Quoted in (ref. 12), 297, 298, 303.

16. Leon S. Moisseiff and Frederick Lienhard. "Suspension Bridges Under the Action of Lateral Forces," with discussion. *Transactions of the American Society of Civil Engineers* 98 (1933), 1080-95, 1096-1141. Quoted in (ref. 12), 299.

17. Henry Petroski. *Design Paradigms: Case Histories of Error and Judgement in Engineering*. New York: Cambridge University Press, 1994.

18. F. B. Farquharson. *Aerodynamic Stability of Suspension Bridges, with Special Reference to the Tacoma Narrows Bridge. Part I: Investigations Prior to October 1941*. Report, Structural Research Laboratory, University of Washington: 1949. Quoted in (ref 17), 157-162.

19. Presidential Commission on the Space Shuttle Challenger Accident. *Report to the President by the Presidential Commission on the Space Shuttle Challenger Accident*, 5 vols. Washington, D.C.: Government Printing Office, 1986.

20. Leonard C. Bruno. "Challenger Explosion." In *When Technology Fails*, edited by Neil Schlager, 609-616. Detroit: Gale Research, 1994

21. Trudy E. Bell and Karl Esch, "The fatal flaw in Flight 51-L." *IEEE Spectrum* 24, no. 2 (1987): 36-51.

22. Diane Vaughan. *The Challenger Launch Decision: Risky Technology, Culture, and Deviance at NASA*. Chicago: University of Chicago Press, 1996.

23. Richard P. Feynman. *"What Do You Care What Other People Think?"* as told to Ralph Leighton. New York: W. W. Norton & Company, 1988.

24. Fredrick P. Brooks, "No Silver Bullet: Essence and Accidents of Software Engineering," IEEE Computer 20, no. 4 (1987): 10-19.

25. Gregory L. Bahnsen. "A Conditional Resolution of the Apparent Paradox of Self-Deception." Ph.D. diss., University of Southern California, 1978.

26. Charles Perrow. *Normal Accidents: Living with High-Risk Technologies*. New York: Basic Books, 1984.

27. Nancy Leveson. *Safeware: System Safety and Computers.* Reading, Mass.: Addison-Wesley Publishing Co., 1995.

28. Mary Shaw. “Prospects for an Engineering Discipline of Software.” Technical report CMU/SEI-90-TR-20. Software Engineering Institute, Carnegie-Mellon University. September, 1980.

29. Nancy Leveson and Clark S. Turner. “An Investigation of the Therac-25 Accidents.” *IEEE Computer* 26, no. 7 (1993): 18-41.

30. J. L. Lions. “Ariane 5 Flight 501 Failure.” Report by the Inquiry Board. Paris: July 1996.

Biography

C. Michael Holloway, NASA Langley Research Center, MS 130/ 1 South Wright Street, Hampton VA 23681-2199, USA, telephone - +1.757.864.1701, facsimile - +1.757.864.4234, e-mail - c.m.holloway@larc.nasa.gov, web – shemesh.larc.nasa.gov/~cmh/

C. Michael Holloway has been a research engineer at the NASA Langley Research Center in Hampton, Virginia since 1983. His interests include accident analysis, programming language theory, development methods for high integrity software, history, theology, and epistemology. He was graduated from the School of Engineering and Applied Science at the University of Virginia with a B.S. in Computer Science in 1983, and did graduate work at the University of Illinois in Champaign-Urbana. He is a member of the IEEE and the IEEE Computer Society. He is married and has two children.